

Fast Exact Max-kernel Search

Ryan Curtin

Parikshit Ram

Alexander G. Gray

Abstract

The wide applicability of kernels makes the problem of max-kernel search ubiquitous and more general than the usual similarity search in metric spaces. We focus on solving this problem efficiently. We begin by characterizing the inherent hardness of the max-kernel search problem with a novel notion of *directional concentration*. Following that, we present a method to use an $O(n \log n)$ algorithm to index any set of objects (points in \mathbb{R}^D or abstract objects) *directly in the Hilbert space* without any explicit feature representations of the objects in this space. We present the first provably $O(\log n)$ algorithm for exact max-kernel search using this index. Empirical results for a variety of data sets as well as abstract objects demonstrate up to 4 orders of magnitude speedup in some cases. Extensions for approximate max-kernel search are also presented.

1 Introduction

Max-kernel search. In this paper, we present a novel algorithm to accelerate the following general task of max-kernel search (MKS): for a given set S of n objects (the *reference set*), a Mercer kernel function $\mathcal{K}(\cdot, \cdot)$, and a query q , find the object $p \in S$ such that:

$$(1) \quad p = \arg \max_{r \in S} \mathcal{K}(q, r).$$

This general form of problem is ubiquitous in computer science; it can be easily seen as a similarity search (for some similarity function $\mathcal{K}(\cdot, \cdot)$). The most simple approach to this general problem is a linear scan over all the objects in S . However, the computational cost of this approach scales linearly with the size of the reference set ($|S| = n$) for a single query, making it computationally prohibitive for large data sets.

A Mercer kernel can define a measure of similarity for classes of objects including points in \mathbb{R}^D , and extending to objects which do not have natural fixed-length representations. Kernels are ubiquitous and can be devised for any new class of objects, such as images and documents (which can be considered as points in \mathbb{R}^D), to more abstract objects like strings (protein sequences [1], text), graphs (molecules [2], brain neuron activation paths), and




Figure 1: Matching images – an example of MKS with the linear, cosine, Gaussian and polynomial kernels as the usual similarity functions.

time series (music, financial data) [3]. The beauty of kernels is the renowned “kernel trick” – the ability to evaluate similarity (inner products) between any pair of objects in some feature space without requiring the explicit feature representations of those objects.

A special case is the problem of nearest-neighbor search (NNS) in metric spaces, in which the closest object to the query with respect to a distance metric is sought. However, the requirement of a distance metric makes many efficient methods for exact and approximate NNS [4] unapplicable to the general MKS.

Some large-scale applications. The widely studied problem of image matching in computer vision is a MKS. The sets of images are of the order of millions and still growing, making linear scan computationally prohibitive. MKS also appears in maximum-a-posteriori inference [5], as well as collaborative filtering (via the widely successful matrix factorization framework) [6]. This matrix factorization obtains an accurate representation of the data in terms of user vectors and item vectors, and the desired result – the user preference of an item – is the inner-product between the two respective vectors (a Mercer kernel). With ever-scaling item sets and millions of users [7], reducing the cost of retrieval of recommendations (which is a MKS) is critical for real-world systems. Finding similar protein/DNA sequences for a query sequence from a large set of biological sequences is also an instance of MKS with biological sequences as the objects and various domain-specific kernels (for example, the p -spectrum

kernel [1], maximal segment match score [8] & Smith-Waterman alignment score [9]¹). An efficient MKS algorithm can be directly applied to biological (sub)sequence matching.

Contributions. To address the breadth of applications of the MKS problem, we present a method to accelerate max-kernel search *for any class of objects with a corresponding Mercer kernel*. Our contributions are:

- The first concept for characterizing the hardness of MKS in terms of the concentration of the kernel values in any direction – the *directional concentration*.
- An $O(n \log n)$ algorithm to index any set of objects *directly in the Hilbert space defined by the kernel* without requiring explicit representations of the objects in this space.
- A novel branch-and-bound algorithm on the index in the Hilbert space, which can achieve orders of magnitude speedups over linear search.
- Value-approximate and order-approximate extensions to the exact max-kernel search algorithm.
- The first $O(\log n)$ runtime bound for *exact* MKS with our proposed algorithm for *any* Mercer kernel.

1.1 Related work

Although there are existing techniques for MKS, *almost all of them solve the approximate problem under restricted settings*. The most common assumption is that the objects are points in some metric space and the kernel function is *shift-invariant* – a monotonic function of the distance between the two objects ($\mathcal{K}(p, q) = f(\|p - q\|)$), such as the Gaussian radial basis function (RBF) kernel. For shift-invariant kernels, a tree-based recursive algorithm has been shown to scale to large sets for maximum-a-posteriori inference [5]. However, a shift-invariant kernel is only applicable to objects already represented in some metric space. In fact, the MKS with a shift-invariant kernel is equivalent to NNS in the input space itself, and can be solved directly using existing methods for NNS, an easier and better-studied problem. For shift-invariant kernels, the points can be explicitly embedded in some low-dimensional metric space such that the inner product between the representations of any two points approximates their corresponding kernel value [10]. This step takes $O(n\mathcal{D}^2)$ time for $S \subset \mathbb{R}^{\mathcal{D}}$ and can be followed by NNS on these representations to accomplish MKS.

Locality-sensitive hashing (LSH) [11] is widely used for image matching, but only with explicitly representable kernel functions which admit a locality sensitive hash-

ing function [12]². Kulis and Grauman [13] apply LSH to solve MKS *approximately* for normalized kernels without any explicit representation. Normalized kernels restrict the self-similarity value to a constant ($\mathcal{K}(x, x) = \mathcal{K}(y, y) \forall x, y \in S$). The preprocessing time for this LSH is $O(p^3)$ and a single query requires $O(p)$ kernel evaluations. Here p controls the accuracy of the algorithm – larger p implies better approximation; the suggest value for p is $O(\sqrt{n})$ with no rigorous approximation guarantees.

A recent work [14] proposed the first technique for exact MKS using a tree-based branch-and-bound algorithm but is restricted only to linear kernels and the algorithm does not have any runtime guarantees. The authors suggest a method for extending their algorithm to non-representable spaces with general Mercer kernels but require $O(n^2)$ preprocessing time.

A case for un-normalized kernels. While some of the kernels used in machine learning (for example, the Gaussian and the cosine kernel) are normalized, some common kernels like the linear kernel (and the polynomial kernels) are not normalized. Many applications require un-normalized kernels: (1) In the retrieval of recommendations, the normalized linear kernel will result in inaccurate user-item preference scores. (2) In biological sequence matching with the domain-specific matching functions, $\mathcal{K}(x, x)$ implicitly corresponds to presence of (genetically) valuable letters (like W, H, P) or invaluable letters (like X)³ in the sequence x . This crucial information is lost in kernel normalization. In this paper, we consider general un-normalized kernels as well as the special case of normalized kernels.

None of the existing techniques can be directly applied to every instance of MKS with general Mercer kernels and any class of objects (Equation 1). Moreover, almost all existing techniques resort to approximate solutions. In this paper, we present a method to perform the exact max-kernel search with $O(n \log n)$ preprocessing time and $O(\log n)$ query time, as well as approximate methods with rigorous accuracy guarantees.

This paper. In the following section, we contrast MKS to NNS, investigate the inherent hardness of MKS and motivate the applicability of indexing schemes used for NNS in MKS. Section 3 discusses the construction of a tree in the kernel space without explicit representations of objects. Section 4 presents the novel branch-and-bound algorithm for exact MKS and extends it to approximate MKS. Sections 5 and 6 examine the theoretical and empirical performance of the approach, respectively. Section 7

²The Gaussian and the cosine kernel admit locality sensitive hashing functions with some modifications.

³See the score matrix for letter pairs in protein sequences at <http://www.ncbi.nlm.nih.gov/Class/FieldGuide/BLOSUM62.txt>.

¹The last two functions provide matching scores for pairs of sequences and can easily be shown to be Mercer kernels.

extends the proposed method to the distributed setting and we conclude with Section 8.

2 Max-kernel search

A Mercer kernel implies that the kernel value for a pair of objects (x, y) corresponds to an inner-product between the vector representation of the objects $\varphi(x), \varphi(y)$ in some inner-product space \mathcal{H} ($\mathcal{K}(x, y) = \langle \varphi(x), \varphi(y) \rangle_{\mathcal{H}}$). Hence every Mercer kernel induces the following metric in \mathcal{H} :

$$(2) \quad d_{\mathcal{K}}(x, y) = \frac{\|\varphi(x) - \varphi(y)\|}{\sqrt{\mathcal{K}(x, x) + \mathcal{K}(y, y) - 2\mathcal{K}(x, y)}}.$$

Whenever MKS can be reduced to NNS in \mathcal{H} , NNS methods for general metrics [4] can be used for efficient MKS. In this section, we show that this reduction is possible only under strict conditions.

Subsequently, we discuss the hardness of MKS and contrast it to the hardness of NNS. Finally, we discuss desirable properties of certain NNS techniques and their applicability to MKS.

Possible reductions and conditions. The nearest neighbor for a query q in \mathcal{H} ($\arg \min_{r \in S} d_{\mathcal{K}}(q, r)$) is the max-kernel candidate (Eq. (1)) if $\mathcal{K}(\cdot, \cdot)$ is a normalized kernel. The two problems can have very different answers with un-normalized kernels. Every kernel also induces a cosine similarity ($\mathcal{K}(x, y) / \sqrt{\mathcal{K}(x, x)\mathcal{K}(y, y)}$) in \mathcal{H} . Similar to the previous case, the object with the maximum cosine-similarity in \mathcal{H} is the max-kernel candidate only for normalized kernels.

Hardness of nearest-neighbor search. However, even if MKS can be reduced to NNS, NNS is still hard for general metrics ($\Omega(n)$ for a single query) without any assumption on the structure of the metric and the set (S, d) . Here we present one notion of characterizing the hardness in terms of the structure of the metric [15]:

Definition 2.1. Let $B_S(p, \Delta) = \{r \in S : d(p, r) \leq \Delta\}$ denote a closed ball of radius Δ around some $p \in S$ with respect to a metric d . Then, the **expansion constant** of (S, d) is defined as the smallest $c \geq 2$ such $|B_S(p, 2\Delta)| \leq c |B_S(p, \Delta)| \quad \forall p \in S$ and $\forall \Delta > 0$.

The expansion constant effectively bounds the number of points that could be sitting on the surface of a hyper-sphere of any radius around any point. If c is high, NNS could be expensive. A value of $c \sim \Omega(n)$ implies that the search cannot be better than linear scan asymptotically. Under the assumption of bounded expansion constant, NNS methods with sub-linear/logarithmic theoretical runtime guarantees have been presented [15, 16].

2.1 Characterizing the hardness of MKS

The kernel values for a query are proportional to the length of the projections in the direction of the query in \mathcal{H} . While the hardness of NNS depends on how concentrated the surface of spheres are (as characterized by the expansion constant), the hardness of MKS should depend on the distribution of the projections in the direction of the query. This distribution can be characterized using the distribution of points in terms of distances. *If two points are close in distance, then their projections in any direction are close as well. However, if two points have close projections in a direction, it is not necessary that the points themselves are close to each other.* It is to characterize this reverse relationship between points (closeness in projections to closeness in distances) that we present a new notion of concentration in any direction:

Definition 2.2. Let $\mathcal{K}(x, y) = \langle \varphi(x), \varphi(y) \rangle_{\mathcal{H}}$ be a Mercer kernel, $d_{\mathcal{K}}(x, y)$ be the induced metric from \mathcal{K} and $B_S(p, \Delta)$ denote the corresponding closed ball of radius Δ around a point p in \mathcal{H} . Let $I_S(v, [a, b]) = \{r \in S : \langle v, \varphi(r) \rangle_{\mathcal{H}} \in [a, b]\}$ be the set of objects in S projected onto a direction v in \mathcal{H} lying in the interval $[a, b]$ along v . Then, the **directional concentration constant** of (S, \mathcal{K}) is defined as the smallest $\gamma \geq 1$ such that $\forall u \in \mathcal{H}$ such that $\|u\|_{\mathcal{H}} = 1$, $\forall p \in S$ and $\forall \Delta > 0$, the set $I_S(u, [\langle u, \varphi(p) \rangle_{\mathcal{H}} - \Delta, \langle u, \varphi(p) \rangle_{\mathcal{H}} + \Delta])$ can be covered by at most γ balls of radius Δ .

The directional concentration constant is not a measure of the number of points projected into a small interval, but rather a measure of the number of “patches” of the data in a particular direction. For a set of points to be close in projections, there are at most γ subsets of points which are close in distances as well. Consider the set of points $B = I_S(q, [a, b])$ projected into an interval in some direction (Figure 2(a)). A high value of γ implies that the number of points in B is high but the points are spread out and the number of balls (with diameter equal to $|b - a|$) required to cover all these points is high as well (each point possibly requiring an individual ball). Figure 2(c) provides one such example. A low value of γ implies that either B has a small size or the size of B is high and B can be covered with a small number of balls (of diameter $|b - a|$). Figure 2(b) is an example of a set with low γ . The directional concentration constant bounds the number of balls of a particular radius required to index points that project onto an interval of size twice the radius.

2.2 Desirable existing techniques

The notion of bounded directional concentration implies that indexing schemes capable of efficiently indexing points in a particular direction might be useful for




Figure 2: Concentration of projections.

the task of MKS. Indexing schemes like space-partitioning trees have been widely used for NNS with success in many cases. Trees offer good hierarchical indexing schemes in low to medium dimensions; for high-dimensional data, trees which exploit some low-dimensional structure have been developed [17, 16]. These hierarchical indexing schemes lend easily to intuitive branch-and-bound algorithms for efficient solutions (especially approximate solutions [18, 19]). In addition, tree-based branch-and-bound algorithms are essentially incremental algorithms, and thus easily extend to anytime algorithms [20]. Most importantly, trees only require a single construction – different algorithms can work on the same tree; in addition, once a tree is built, point insertions and deletions are generally cheap.

Which tree to use for MKS? Given the numerous advantages of trees, we wish to select an appropriate tree for efficient MKS. The following two properties are desired of any indexing scheme used for MKS:

- Explicit representation of the objects is not required.
- It should have sub-quadratic construction time.

The *kd*-tree [21] and the metric tree [22] exhibit good characteristics and are widely used in NNS. However, the *kd*-tree requires the explicit representation of the points $\varphi(x)$ in \mathcal{H} for its axis-aligned splits. For similar reasons, random-projection trees [17] and PCA-trees [23] cannot be used for MKS. Metric trees can be constructed without the explicit representations since they can work with only the ability to evaluate the induced metric d_K [24]⁴. In this paper, we consider the recently formulated *cover tree* [16] for MKS. It provides a systematic way to build a tree without explicit object representations and with a sub-quadratic construction time. A key difference between cover trees and the aforementioned trees is that the *kd*-tree and the metric tree are binary trees while the cover trees can have multiple number of children. Moreover,

⁴Calculation of the explicit mean μ of a node is avoidable by using the kernel trick for operations on the mean: $\langle \mu, \varphi(q) \rangle = \sum_{x \in T} \mathcal{K}(x, q) / |T|$ (where T is the set of points in the node). However, this makes the tree construction and tree search computational prohibitive for efficient MKS.

the time complexities of building and querying a cover trees have been analyzed extensively [16, 25], whereas *kd*-trees and metric trees have been analyzed only under very limited settings [21].

3 Indexing in kernel space with cover trees

A cover tree stores a dataset S of size n in the form of a levelled tree. The structure has $O(n)$ space requirement (Theorem 1 in [16]). Each level is a “cover” for the level beneath it and is indexed by an integer scale i which decreases as the tree is descended. Let C_i denote the set of nodes at scale i . For all scales i , the following invariants hold: (i) (nesting invariant) $C_i \subset C_{i-1}$, (ii) (covering tree invariant) $\forall p \in C_{i-1}, \exists q \in C_i: d(p, q) \leq 2^i$, and exactly one such q is a parent of p . (iii) (separation invariant) For all $p, q \in C_i$, $d(p, q) > 2^i$.

Although the cover tree is defined as infinite levelled sets, the tree has a precise finite representation. As explained in [25], “The *implicit representation* consists of infinitely many levels C_i with the level C_∞ containing a single node which is the root and the level $C_{-\infty}$ containing every point in the dataset as a node. The *explicit representation* is required to store the tree in $O(n)$ space. It coalesces all nodes in the tree for which the only child is the self-child. This implies that every explicit node either has a parent other than the self-parent or has a child other than a self-child.”

Tree construction in the kernel space \mathcal{H} . Every node in a cover tree is associated with a single point p . Hence, the cover tree construction only requires distances between the points in the set (the tree construction algorithm is presented in Section A of the supplement); and therefore construction in \mathcal{H} *does not require any explicit representation in \mathcal{H}* . From Equation 2, the pairwise distance $d_K(x, y)$ can be evaluated with only three kernel evaluations. If the self-kernel values $\mathcal{K}(x, x) \forall x \in S$ are precomputed and saved, $d_K(x, y)$ only requires one evaluation of $\mathcal{K}(x, y)$. The tree construction time is bounded by the following theorem:

Theorem 3.1. (Theorem 3.6 in [26]) For a data set S of n objects and a metric d_K with expansion constant c , the tree construction requires at most $O(c^6 n \log n)$ time.

Remark. We would like to point out that while we consider the cover tree to implicitly index points in \mathcal{H} without any modification, the cover tree has only been used for NNS to this date. The novelty of this paper is a new branch-and-bound algorithm using this tree (Section 4) to solve the general task of MKS with provable theoretical guarantees (Section 5) and supporting empirical evidence (Section 6).

4 Simple branch-and-bound algorithm

In this section, we present a simple branch-and-bound algorithm on the cover tree for MKS. Branch-and-bound is widely used in NNS with the help of the triangle inequality of the distance metric. In the absence of the triangle inequality for kernel functions, we obtain a novel bound on the max-kernel value possible between a query and any subtree of a cover tree. Then we present the algorithm for exact and approximate search.

Bounding the max-kernel value. A cover tree node is defined by an object p and a level i . Let S_p^i denote the set of objects in the subtree rooted at a node defined by object p at level i . In the following theorem, we bound the maximum possible kernel value between a query and an object in the subtree of the cover tree. For notational convenience, we will denote $\max_{r \in R} K(q, r)$ as $K(q, R)$.

Theorem 4.1. Given a cover tree node rooted at an object p at level i in the kernel space \mathcal{H} and a (query) object q , the maximum kernel function value between q and any object in the set S_p^i is bounded as:

$$(3) \quad K(q, S_p^i) \leq K(q, p) + 2^{i+1} \sqrt{K(q, q)}.$$

Proof. Suppose that p^* is the best possible match in the set S_p^i for q and let \vec{u} be a unit vector in the direction of the line joining $\varphi(p)$ to $\varphi(p^*)$ in \mathcal{H} . Then $\varphi(p^*) = \varphi(p) + \Delta \cdot \vec{u}$ where $\Delta = d_K(p, p^*)$ is the distance $\varphi(p)$ and the best possible match $\varphi(p^*)$ (see Figure 3). Then we have the following:

$$\begin{aligned} K(q, S_p^i) &= K(q, p^*) = \langle \varphi(q), \varphi(p^*) \rangle_{\mathcal{H}} \\ &= \langle \varphi(q), \varphi(p) + \Delta \cdot \vec{u} \rangle_{\mathcal{H}} \\ (4) \quad &\leq \langle \varphi(q), \varphi(p) \rangle_{\mathcal{H}} + \Delta \|\varphi(q)\|_{\mathcal{H}}, \end{aligned}$$

where the last inequality follows from the Cauchy-Schwartz inequality ($\langle x, y \rangle \leq \|x\| \|y\|$) and the fact that $\|\vec{u}\| = 1$. From the definition of the kernel function, Equation 4 gives us $K(q, S_p^i) \leq K(q, p) + \Delta \sqrt{K(q, q)}$. We

bound Δ from above using the covering invariant – for any cover tree node p at level i , the distance to the farthest child node is bounded by 2^i . Applying this bound recursively with the triangle inequality of d_K gives us $\Delta = d_K(p, p^*) \leq \sum_{j=-\infty}^i 2^j = 2^{i+1}$. The statement of the theorem follows. \square




Figure 3: Max-kernel upper bound.

For normalized kernels ($K(x, x) = 1 \forall x \in S$), all the points are on the surface of a hyper-sphere in \mathcal{H} . In this case, the above bound in Theorem 4.1 is correct but possibly loose. In the following theorem, we present a tighter bound specifically for this condition:

Theorem 4.2. Consider a kernel K such that $K(x, x) = 1 \forall x \in S$. Given a cover tree node rooted at an object p at level i in \mathcal{H} and a (query) object q , the maximum kernel function value between q and any object in the set S_p^i is bounded as:

$$(5) \quad K(q, S_p^i) \leq \begin{cases} K(q, p)(1 - 2^{2i+1}) \\ \quad + 2^{i+1} \sqrt{(1 - K(q, p)^2)(1 - 2^{2i})}, \\ \quad \text{if } K(q, p) \leq 1 - 2^{2i+1} \\ 1.0, \text{ otherwise} \end{cases}$$

The proof is similar to the proof of Theorem 4.1 and presented in Section B of the supplement.

Algorithm 1 FastMKS(Tree T , query q)

Initialize $R_{\infty} = C_{\infty}$ // the root of the tree
for $i = \infty$ **to** $-\infty$ **do**
 $R = \{\text{Children}(r) : r \in R_i\}$ // tree descend
 $R_{i-1} = \{r \in R : K(q, r) \geq K(q, R) - 2^i \sqrt{K(q, q)}\}$
end for
return $\arg \max_{r \in R_{-\infty}} K(q, r)$ // the leaf nodes.

The branch-and-bound algorithm. The bound on $K(q, S_p^i)$ is used to decide whether a node is retained for further exploration or removed (“pruned”) from consideration. The branch-and-bound algorithm **FastMKS**(T, q) is presented in Alg. 1. If the maximum possible kernel value from a subtree is less than the current best kernel value, that subtree is not explored any further. The




Figure 4: Branch-and-bound tree-traversal – the green nodes are retained and the red nodes are pruned.

best possible kernel value from a subtree and a step of the algorithm is depicted in Figure 4. For a retained node, all its children are explored. The correctness of **FastMKS**(T, q) follows from:

Theorem 4.3. *If T is a cover tree on S , then **FastMKS**(T, q) (Alg. 1) returns $\arg \max_{r \in S} \mathcal{K}(q, r)$.*

We present a sketch proof. The complete proof is presented in Section C of the supplement.

Proof. (sketch) By Theorem 4.1, $\mathcal{K}(q, R_{i-1}) \geq \mathcal{K}(q, S) - 2^i \sqrt{\mathcal{K}(q, q)} \forall i$. Therefore, $\lim_{i \rightarrow -\infty} \mathcal{K}(q, R_{i-1}) = \mathcal{K}(q, S)$. Hence $\arg \max_{r \in R_{-\infty}} \mathcal{K}(q, r) = \arg \max_{r \in S} \mathcal{K}(q, r)$. \square

Approximate max-kernel search. Similarity search problems can be approximated for further scalability. Even though we are focusing on exact max-kernel search in this paper, we wish to demonstrate that the tree based method can be very easily extended to perform the approximate max-kernel search. Approximation can be achieved in the following ways:

1. Absolute value approximation (AVA): return $p \in S$ such that $\mathcal{K}(q, p) \geq \mathcal{K}(q, S) - \epsilon$.
2. Relative value approximation (RVA): return $p \in S$ such that $\mathcal{K}(q, p) \geq \mathcal{K}(q, S) - \epsilon |\mathcal{K}(q, S)|$.
3. Rank approximation (RA): return $p \in S$ such that $|\{r \in S : \mathcal{K}(q, r) > \mathcal{K}(q, p)\}| \leq \tau$.

Care has to be taken for relative value approximation since there is no guarantee that $\mathcal{K}(q, S) > 0$. The following stopping rules can be used at the end of an iteration in the for loop in **FastMKS**(T, q) (Alg. 1) for the value approximations. The best candidate up until then is the approximate solution.

1. AVA stopping: if $\epsilon \geq 2^i \sqrt{\mathcal{K}(q, q)}$, stop.
2. RVA stopping: assuming $\mathcal{K}(q, S) > 0$, if $\mathcal{K}(q, R_{i-1}) \geq (2^i/\epsilon) \sqrt{\mathcal{K}(q, q)}$, stop.⁵

Theorem 4.4. *The AVA stopping rule returns a point $p \in S$ such that $\mathcal{K}(q, p) \geq \mathcal{K}(q, S) - \epsilon$.*

Proof. At the end of iteration at level i , $\mathcal{K}(q, R_{i-1}) \geq \mathcal{K}(q, S) - 2^i \sqrt{\mathcal{K}(q, q)}$. If $2^i \sqrt{\mathcal{K}(q, q)} \leq \epsilon$, the AVA condition is satisfied. \square

⁵The stopping rule can be easily modified for $\mathcal{K}(q, S) < 0$.

Theorem 4.5. *Assuming $\mathcal{K}(q, S) > 0$, the RVA stopping rule returns a point $p \in S$ such that $\mathcal{K}(q, p) \geq (1 - \epsilon) \mathcal{K}(q, S)$.*

Proof. When $\mathcal{K}(q, R_{i-1}) \geq (2^i/\epsilon) \sqrt{\mathcal{K}(q, q)}$, then $2^i \sqrt{\mathcal{K}(q, q)} \leq \epsilon \mathcal{K}(q, R_{i-1}) \leq \epsilon \mathcal{K}(q, S)$. Using this, we have $\mathcal{K}(q, R_{i-1}) \geq \mathcal{K}(q, S) - 2^i \sqrt{\mathcal{K}(q, q)} \geq \mathcal{K}(q, S) - \epsilon \mathcal{K}(q, S)$. \square

Rank-approximation can be achieved by performing stratified sampling on the cover tree [19]. The technique is more involved and presented in Section D of the supplement for the lack of space.

5 Runtime analysis

The runtime analysis of **FastMKS** will make use of the following results [16]:

Lemma 5.1. *The number of children of any node p is bounded by c^4 .*

Lemma 5.2. *The maximum depth of any point p in the explicit representation is $O(c^2 \log n)$.*

The main result of this section is the search time complexity of **FastMKS**(T, q) in terms of the number of objects in S and the properties of (S, \mathcal{K}) :

Theorem 5.1. *Given a Mercer kernel \mathcal{K} , if the dataset S of size n has an expansion constant c (with the metric $d_{\mathcal{K}}$) and a directional concentration constant γ , **FastMKS**(T, q) requires $O(c^{12} \gamma^2 \log n)$ time.*

Proof. The first part of the proof is similar to the runtime analysis of NNS with cover trees [16]. Let R^* denote the last explicit R_i considered by the algorithm. By Lemma 5.2, the explicit depth of any point in R^* is at most $k = O(c^2 \log n)$. The maximum number of iterations required would be at most $k |R^*| \leq k \max_i |R_i|$. The amount of work done in each iteration is at most $O(\max_i |R_i|)$, hence resulting in a total of $O(k \max_i |R_i|^2)$ work. Moreover, from Lemma 5.1, the total number of children encountered throughout the whole algorithm is at most $k \max_i |R_i| \cdot c^4$. Hence, the pruning/retaining rule does at most $O(k \max_i |R_i| \cdot c^4)$ work. Also, $R_{-\infty} \leq \max_i |R_i|$. Therefore, the algorithm requires at most $O(k \max_i |R_i|^2 + k \max_i |R_i| \cdot c^4)$ time.

Now we bound $\max_i |R_i|$. Let $u = \varphi(q)/\|\varphi(q)\|$. Then $I_S(\varphi(q), [a, b]) = I_S(u, [a/\|\varphi(q)\|, b/\|\varphi(q)\|])$. For any level i , let $R = \{\text{Children}(r) : r \in R_i\}$ and let $\kappa = \mathcal{K}(q, R)$ and $\kappa^* = \mathcal{K}(q, S)$. Then

$$\begin{aligned} R_{i-1} &= \{r \in R : \mathcal{K}(q, r) \geq \kappa - 2^i \|\varphi(q)\|\} \\ &= I_S(\varphi(q), [\kappa - 2^i \|\varphi(q)\|, \kappa^*]) \cap R \\ &\subseteq I_S(\varphi(q), [\kappa - 2^i \|\varphi(q)\|, \kappa^*]) \cap C_{i-1} \\ &\subseteq I_S(\varphi(q), [\kappa^* - 2^{i+1} \|\varphi(q)\|, \kappa^*]) \cap C_{i-1} \end{aligned}$$

since $\kappa^* \leq \kappa + 2^i \|\varphi(q)\|$.

Then for any $r \in I_S(\varphi(q), [\kappa^* - 2^{i+1} \|\varphi(q)\|, \kappa^*])$,

$$\begin{aligned} I_S(\varphi(q), [\kappa^* - 2^{i+1} \|\varphi(q)\|, \kappa^*]) \\ \subseteq I_S(q, [\mathcal{K}(q, r) - 2^{i+1} \|\varphi(q)\|, \mathcal{K}(q, r) + 2^{i+1} \|\varphi(q)\|]) \\ = I_S(u, [\langle u, \varphi(r) \rangle - 2^{i+1}, \langle u, \varphi(r) \rangle + 2^{i+1}]). \end{aligned}$$

By the definition of the directional concentration constant, there exists r_j s such that:

$$\begin{aligned} I_S(u, [\langle u, \varphi(r) \rangle - 2^{i+1}, \langle u, \varphi(r) \rangle + 2^{i+1}]) \\ \subseteq \cup_{j=1}^{\gamma} B_S(r_j, 2^{i+1}). \end{aligned}$$

Bounding the size of $I_S(q, [\mathcal{K}(q, r) - 2^{i+1} \|\varphi(q)\|, \mathcal{K}(q, r) + 2^{i+1} \|\varphi(q)\|]) \cap C_{i-1}$ amounts to bounding the number of disjoint balls of radius 2^{i-2} that can be packed into each of the γ balls $B_S(r_j, 2^{i+1} + 2^{i-2})$. For each of the r_j , we have:

$$\begin{aligned} |B_S(r_j, 2^{i+1} + 2^{i-2})| &\leq |B_S(r', 2^{i+2} + 2^{i-1})| \\ &\leq |B_S(r', 2^{i+3})| \leq c^5 |B_S(r, 2^{i-2})|. \end{aligned}$$

Hence, $\forall i, |R_i| \leq \gamma c^5$. Hence $\max_i |R_i| \leq \gamma c^5$, thus giving us the statement of the theorem. \square

Comparing to the query time $O(c^{12} \log n)$ for NNS [16], it is clear that **FastMKS** has similar $\log n$ scaling, but also has an extra price of γ^2 for solving the more general problem of max-kernel search.

6 Evaluation

We evaluate the **FastMKS** algorithm with different kernels and datasets. For every experiment, we query the top $\{1, 2, 5, 10\}$ max-kernel candidates and report the speedup over linear search (in terms of the number of kernel evaluations performed). The cover tree and the algorithm is developed in MLPACK [27] with the implementation details in Section E of the supplement.⁶

Datasets. We use two different classes of datasets – (1) Datasets with fixed-length objects: These include the MNIST dataset [28], the Isomap “Images” dataset,

several datasets from the UCI machine learning repository [29], three collaborative filtering datasets (MovieLens, Netflix [30], Yahoo! Music [7]), the LCDM astronomy dataset [31], the LiveJournal blog moods text dataset [32] and a subset of the 80 million tiny images dataset [33] (the sizes of the datasets are presented in Table 6). (2) Datasets without fixed length representation: We consider protein sequences from the GenBank⁷.

Datasets	$ Q $	$ S $	dims
Y! Music	100000	600000	50
MovieLens	6040	3706	50
Opt-digits	450	1347	64
Physics	37500	112500	78
Bio	75000	210409	74
Covertypes	150000	431012	55
LiveJournal	10000	10000	25327
MNIST	10000	60000	784
Netflix	480189	17770	50
Corel	10000	27749	32
LCDM	6000000	10777216	3
TinyImages	5000	1000000	384

Table 1: Details of the vector datasets.

Kernels. We consider all of the following kernels for the vector datasets: *cosine*, *hyperbolic tangent*⁸, *linear*, and *polynomial* (with degree 10). While **FastMKS** is applicable to any kernel, MKS with the Gaussian kernel reduces to NNS in the input space; hence, we omit this kernel from our experiments. The p -spectrum kernel [1] is used for the protein sequence data.

Results. The results for the vector datasets are summarized in Figure 5. While the speedups range from anywhere between 1 to over 10^4 , a speedup close to an order of magnitude is seen in most large datasets (except MNIST). It is also quite clear that different kernels give very different speedups for the same dataset. This can be attributed to the fact that the expansion constant and the directional concentration constant are properties of the dataset-kernel pair. The results for the protein sequence data (Figure 6) indicate that the speedups increase with increasing reference set size, recording a speedup of over two orders of magnitude for a set of around 100000 sequences, exhibiting the logarithmic scaling of **FastMKS**.


7 Distributed data

Despite the theoretical and empirical efficiency of **FastMKS**, the algorithm still requires the dataset (and


⁷<ftp://ftp.ncbi.nih.gov/refseq/release/complete/>

⁸While the hyperbolic tangent kernel is not Mercer in general, our proposed algorithm works correctly with non-Mercer kernels if the kernel is positive definite when restricted to the dataset.

⁶See <http://www.mlpac.org/> for more information on MLPACK.



(a) Cosine and hyperbolic tangent kernels



(b) Linear and polynomial (degree 10) kernels

Figure 5: Speedups over linear scan for **FastMKS** with $k = 1, 2, 5, 10$.




Figure 6: Speedups over linear scan for **FastMKS** with $k = 1, 2, 5, 10$ with sets of protein sequences of increasing size to demonstrate the scaling of **FastMKS**.

the tree) to fit completely in the memory of a single machine. Large datasets generally need to be distributed across multiple nodes to load the whole data in memory. While a tree can be constructed on distributed data, this tree is inefficient for MKS because of the high internode communication in the branch-and-bound algorithm. In this section, we discuss a method to extend **FastMKS** to the distributed data setting and compare it to the obvious baseline. Since communication and data exchange are major bottlenecks in parallel and distributed systems, our proposed method and the baseline avoid these bottlenecks as well as fit into the map-reduce framework.

Linear scan baseline. Indexing schemes are generally not conducive to distributed data and distributed linear scan is a possible alternative. Let the dataset of size n be evenly distributed over m nodes with each node containing (n/m) objects and let there be a master node that communicates with each of these m nodes. For single query, the master sends the query to each of the m nodes

in the map phase and the mappers on every node perform a linear scan in parallel. The best match from each of these nodes is returned to the master node which can then return the best among those m returned matches in the reduce phase. Then the total time required to obtain the best match is $(c_1 m + c_2 \frac{n}{m})$. It is important to note that each of the m nodes containing the actual data performs (n/m) kernel evaluations while the master node just sorts m kernel values; therefore, $c_1 \ll c_2$. The minimum possible runtime achievable by linear scan in this distributed system is

$$2\sqrt{c_1 c_2 n} \sim O(\sqrt{n}) \text{ with } m = \sqrt{c_2 n / c_1}.$$

Multiple tree indices. Let us consider this following scheme for indexing: for each of the m nodes containing (n/m) objects, build and save a cover tree on the data in each of the nodes. Similar to the linear scan case, for a single query, the master sends the query to each of the m nodes in the map phase and the mappers on every node perform the branch-and-bound search on the cover trees in parallel. The best matches from each of m trees are returned to the master node which then reports the best among them in the reduce phase. Since we have shown that a single query on a cover tree requires logarithmic time, let the time taken for this parallel search on multiple trees be $(c_1 m + c'_2 \log \frac{n}{m})$, where c'_2 is the scaling constant independent of n for the cover tree on this dataset. In this scenario, the minimum possible achievable runtime is

$$c'_2 \log(e n c_1 / c'_2) \sim O(\log n) \text{ with } m = c'_2 / c_1.$$

Remark. The optimal values of m can only be used if there are enough number of nodes available and if (n/m) is small enough for the data in the node to fit in memory. However, searching over multiple tree indices is a simple yet efficient extension of our proposed method for distributed data. In the best case, our method can achieve

$O(\log n)$ scaling while the distributed linear scan can only achieve $O(\sqrt{n})$ scaling. More non-trivial extensions using recent techniques [34, 35] can achieve further efficiency in the distributed setting.

8 Conclusion

Our work appears to be the most comprehensive study of the general MKS problem to date, including the first rigorous characterization of its hardness and the first general-kernel method with provably logarithmic query time, sub-quadratic preprocessing time, and speed in practice. A tighter theoretical analysis without the bounded directional concentration constant assumption and an empirical investigation on more abstract objects like graphs and time series will be part of our future work. We wish to extend our search algorithm to more classes of similarity functions beyond Mercer kernels.

Appendix

A Cover tree construction in Hilbert space

The details of the cover tree construction are provided in Algorithm 2. For a cover tree built on dataset S , each node is only associated with a single point $p \in S$. Therefore, because the only distance computations involve points in S , the explicit representation of the objects in \mathcal{H} is not required (by the kernel trick).

Given a Mercer kernel \mathcal{K} for a class of objects, the distances between points in \mathcal{H} required for the tree construction in \mathcal{H} can be evaluated using the distance metric $d_{\mathcal{K}}$ induced from the kernel. Three kernel evaluations are required to compute the distance; however, if the self-kernel values $\mathcal{K}(x, x) \forall x \in S$ are precomputed and saved, $d_{\mathcal{K}}(x, y)$ can be evaluated with a single evaluation of $\mathcal{K}(x, y)$.

Algorithm 2 **Construct**($p, \langle \text{NEAR}, \text{FAR} \rangle, i$) [26]

```

if NEAR =  $\emptyset$  then
    return  $\{p, \emptyset\}$ .
else
    {SELF, NEAR} =
        Construct( $p, \text{Split}(d(p, \cdot), 2^{i-1}, \{\text{NEAR}\}), i - 1$ )
    add SELF to Children( $p$ )
    while NEAR  $\neq \emptyset$  do
        pick  $q$  in NEAR
        {CHILD, UNUSED} =
            Construct( $q, \text{Split}(d(q, \cdot), 2^{i-1}, \{\text{NEAR}, \text{FAR}\}),$ 
                 $i - 1$ )
        add CHILD to Children( $p$ )
        {NEW-NEAR, NEW-FAR} =
            Split( $d(p, \cdot), 2^i, \{\text{UNUSED}\}$ )
        FAR  $\leftarrow$  FAR  $\cup$  NEW-FAR
        NEAR  $\leftarrow$  NEAR  $\cup$  NEW-NEAR
    end while
    return  $\{p, \text{FAR}\}$ .
end if

Split( $d(p, \cdot), r, \{S_1, S_2, \dots\}$ )
    NEAR =  $\bigcup_i \{q \in S_i : d(p, q) \leq r\}$ 
    FAR =  $\bigcup_i \{q \in S_i : 2r > d(p, q) > r\}$ 
     $\forall i, S_i \leftarrow S_i \setminus (\text{NEAR} \cup \text{FAR})$ .
    return {NEAR, FAR}.

```

The algorithm calls the recursive function **Construct**($p, \langle \text{NEAR}, \text{FAR} \rangle, i$) where p is a point, $\langle \text{NEAR}, \text{FAR} \rangle$ are the point sets and i is the current level of the tree. This recursive function calls a subroutine **Split**($d(p, \cdot), r, \{S_1, S_2, \dots\}$) with $d(p, \cdot)$ representing the set of distances of the points in the point sets $\{S_1, S_2, \dots\}$ to the point p and r is the splitting distance to form the NEAR and FAR sets.

B Proof of Theorem 4.2

We restate the theorem for completeness:

Theorem B.1. *Consider a kernel \mathcal{K} such that $\mathcal{K}(x, x) = 1 \forall x \in S$. Given a cover tree node rooted at an object p at level i in \mathcal{H} and a (query) object q , the maximum kernel function value between q and any object in the set S_p^i is bounded as:*

$$(6) \quad \mathcal{K}(q, S_p^i) \leq \begin{cases} \mathcal{K}(q, p)(1 - 2^{2i+1}) \\ \quad + 2^{i+1} \sqrt{(1 - \mathcal{K}(q, p)^2)(1 - 2^{2i})}, \\ \quad \text{if } \mathcal{K}(q, p) \leq 1 - 2^{2i+1} \\ 1.0, \text{ otherwise} \end{cases}$$

Proof. Since all the points are sitting on the surface of a hypersphere in \mathcal{H} , $\mathcal{K}(q, p)$ denotes the cosine of the angle made by $\varphi(q)$ and $\varphi(p)$ at the origin. If we first consider the case where q lies within the ball bounding cover tree node p at level i (that is, if $d_K(q, p) < 2^{i+1}$), it is clear that the maximum possible kernel evaluation should be 1, because there could exist a point in S_p^i whose angle to q is 0. We can easily modify this condition to an easier condition on $\mathcal{K}(q, p)$, which is $\mathcal{K}(q, p) < 1 - 2^{2i+1}$.

Now, for the other case, let $\cos \theta_{qp} = \mathcal{K}(q, p)$ and $p^* = \arg \max_{r \in S_p^i} \mathcal{K}(q, r)$. Let θ_{pp^*} be the angle between $\varphi(p)$ and $\varphi(p^*)$ and θ_{qp^*} be the angle between $\varphi(q)$ and $\varphi(p^*)$ at the origin. Then

$$(7) \quad \mathcal{K}(q, p^*) = \cos \theta_{qp^*} \leq \cos(\{\theta_{qp} - \theta_{pp^*}\}_+) .$$

Now we know that $d_K(p, p^*) \leq 2^{i+1}$ and that $d_K(p, p^*) = \sqrt{2 - 2 \cos \theta_{pp^*}}$. Hence $\cos \theta_{pp^*} \geq 1 - 2^{2i+1}$, and $\theta_{pp^*} \leq |\cos^{-1}(1 - 2^{2i+1})|$ and take $\theta = |\cos^{-1}(1 - 2^{2i+1})|$. Combining this with equation (7), we get:

$$\mathcal{K}(q, S_p^i) \leq \cos(\{\theta_{qp} - \theta_{pp^*}\}_+) \leq \cos(\{\theta_{qp} - \theta\}_+) .$$

Substituting the value of θ above and simplifying gives us the statement of the theorem. \square

C Proof of Theorem 4.3

We will restate the theorem here:

Theorem C.1. *If T is a cover tree on S , then $\text{FastMKS}(T, q)$ returns $\arg \max_{r \in S} \mathcal{K}(q, r)$.*

Proof. Let $p^* = \arg \max_{r \in S} \mathcal{K}(q, r)$ and $\kappa^* = \mathcal{K}(q, p^*)$. At the beginning of the iteration at level ∞ , p^* is in consideration since $p^* \in S$ and for $p \in C_\infty$, $S_p^\infty = S$. At the end of the iteration at level i , all points $p \in S$ such that $\mathcal{K}(q, p) \geq \mathcal{K}(q, R_{i-1}) - 2^i \sqrt{\mathcal{K}(q, q)}$ are still in consideration. Since $\kappa^* \geq \mathcal{K}(q, R_{i-1})$, p^* is still in consideration. Either $p^* \in R_{i-1}$ or p^* is the

grandchild of some point $p \in R_{i-1}$. Hence, by theorem 4.1, $\mathcal{K}(q, R_{i-1}) \geq \kappa^* - 2^i \sqrt{\mathcal{K}(q, q)} \forall i$. Now $\lim_{i \rightarrow -\infty} \mathcal{K}(q, R_{i-1}) \geq \lim_{i \rightarrow -\infty} (\kappa^* - 2^i \sqrt{\mathcal{K}(q, q)}) = \kappa^*$ (assuming $\mathcal{K}(q, q) < \infty$). Hence $\mathcal{K}(q, R_{-\infty}) = \kappa^*$ and $\text{FastMKS}(T, q)$ returns $\arg \max_{r \in R_{-\infty}} \mathcal{K}(q, r) = p^*$. \square

D Rank-approximate max-kernel search

The rank-approximation of the max-kernel search is defined as follows: for a given set S of n objects (the reference set), a (Mercer) kernel function $\mathcal{K}(\cdot, \cdot)$, and a query q , find the object $p \in S$ such that

$$|\{r \in S: \mathcal{K}(q, r) > \mathcal{K}(q, p)\}| \leq \tau .$$

As mentioned in Ram et.al, 2009 [19], the idea is to draw enough samples S' from the tree such that

$$\Pr(|\{r \in S: \mathcal{K}(q, r) > \mathcal{K}(q, S')\}| < \tau) \geq 1 - \delta .$$

Simplifying the formulation presented in [19], the probability of always missing the top τ values for a given query q after k samples with replacement is given by $(1 - \frac{\tau}{n})^k$. If we want a $(1 - \delta)$ success rate of sampling, then we want k to be such that

$$\left(1 - \frac{\tau}{n}\right)^k < \delta, \quad \text{and} \quad \left(1 - \frac{\tau}{n}\right)^{k-1} > \delta,$$

giving $k = \left\lceil \frac{\log \delta}{\log(1 - \frac{\tau}{n})} \right\rceil$. Given that k samples (as defined above) is to be made from the tree, the stratified sampling on a tree is presented in Algorithm 3. We assume that at each node of the tree, we have access to the number of points in the subtree S_r^i rooted at node r at level i for every node in the tree. This algorithm returns a τ -rank approximate solution to the max-kernel operation with probability $(1 - \delta)$.

Algorithm 3 RAFastMKS(Tree T , query q , rank error τ , failure probability δ)

```

Initialize  $R_\infty = C_\infty$  // the root of the tree
Set  $k = \left\lceil \frac{\log \delta}{\log(1 - \frac{\tau}{n})} \right\rceil$  // the number of samples
for  $i = \infty$  to  $-\infty$  do
     $R = \{\text{Children}(r) : r \in R_i\}$  // tree descend
     $R' = \{r \in R : \mathcal{K}(q, r) \geq \mathcal{K}(q, R) - 2^i \sqrt{\mathcal{K}(q, q)}\}$ 
     $R_{i-1} = \{r \in R' : |S_r^i| > \frac{n}{k}\}$  // approximate by sampling
     $R_{-\infty} = R_{-\infty} \cup \{R' \setminus R_{i-1}\}$ 
end for
return  $\arg \max_{r \in R_{-\infty}} \mathcal{K}(q, r)$  // the leaf nodes.
```

E Implementation details

We use the following performance-improving optimizations on the cover tree:

- (1) Instead of the upperbound of 2^{i+1} for the distance of a cover tree node p at level i to its furthest descendant, the actual distance to the furthest descendant is cached at tree construction time and used at search time,
- (2) The distance to the parent is cached to avoid evaluating \mathcal{K} for the child nodes where an improvement is impossible,
- (3) Instead of the base 2, an experimentally verified base of 1.3 is used for best results [26],
- (4) $\mathcal{K}(x, x) \forall x \in S$ is precomputed and stored for future use to reduce the number of kernel evaluations required for a single distance computation $d_{\mathcal{K}}(x, y)$ to one (just $\mathcal{K}(x, y)$ since $\mathcal{K}(x, x)$ and $\mathcal{K}(y, y)$ are pre-computed),
- (5) We use the tighter bound in Theorem 4.2 for normalized kernels.

References

- [1] C. Leslie, E. Eskin, and W.S. Noble. The Spectrum Kernel: A String Kernel for SVM Protein Classification. In *Proc. of the Pacific Symp. on Biocomputing*, 2002.
- [2] K.M. Borgwardt, C.S. Ong, S. Schönauer, S. V. N. Vishwanathan, A.J. Smola, and H.P. Kriegel. Protein Function Prediction via Graph Kernels. *Bioinformatics*, 2005.
- [3] K. Müller, A. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V. Vapnik. Predicting Time Series with Support Vector Machines. *Intl. Conf. on Artificial Neural Networks*, 1997.
- [4] K.L. Clarkson. Nearest-neighbor Searching and Metric Space Dimensions. *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, 2006.
- [5] M. Klaas, D. Lang, and N. de Freitas. Fast Maximum-a-posteriori Inference in Monte Carlo State Spaces. In *Artificial Intelligence and Statistics*, 2005.
- [6] Y. Koren, R. M. Bell, and C. Volinsky. Matrix Factorization Techniques for Recommender Systems. *IEEE Computer*, 2009.
- [7] G. Dror, N. Koenigstein, Y. Koren, and M. Weimer. The Yahoo! Music Dataset and KDD-Cup’11. *Journal Of Machine Learning Research*, 2011.
- [8] S.F. Altschul, W. Gish, W. Miller, E.W. Myers, D.J. Lipman, et al. Basic Local Alignment Search Tool. *Journal of Molecular Biology*, 1990.
- [9] T. F. Smith and M. S. Waterman. Identification of Common Molecular Subsequences. *Journal of Molecular Biology*, 1981.
- [10] A. Rahimi and B. Recht. Random Features for Large-scale Kernel Machines. *Advances in Neural Info. Proc. Systems 20*, 2007.
- [11] A. Gionis, P. Indyk, and R. Motwani. Similarity Search in High Dimensions via Hashing. *Proceedings of the 25th Intl. Conf. on Very Large Data Bases*, 1999.
- [12] M. S. Charikar. Similarity Estimation Techniques from Rounding Algorithms. In *Proceedings of the 34th annual ACM Symp. on Theory of Comp.*, 2002.
- [13] B. Kulis and K. Grauman. Kernelized Locality-sensitive Hashing for Scalable Image Search. In *IEEE 12th Intl. Conf. on Computer Vision*, 2009.
- [14] P. Ram and A.G. Gray. Maximum Inner-Product Search Using Cone Trees. In *SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2012.
- [15] D. R. Karger and M. Ruhl. Finding Nearest Neighbors in Growth-Restricted Metrics. *Proc. of ACM Symp. on Theory of Computing*, 2002.
- [16] A. Beygelzimer, S. Kakade, and J.C. Langford. Cover Trees for Nearest Neighbor. *Proc. of the Intl. Conf. on Machine Learning*, 2006.
- [17] S. Dasgupta and Y. Freund. Random projection trees and low dimensional manifolds. In *Proc. of ACM Symp. on Theory of Comp.*, 2008.
- [18] P. Ciaccia and M. Patella. PAC Nearest Neighbor Queries: Approximate and Controlled Search in High-dimensional and Metric spaces. *Proceedings of 16th Intl. Conf. on Data Engineering*, 2000.
- [19] P. Ram, D. Lee, H. Ouyang, and A. G. Gray. Rank-Approximate Nearest Neighbor Search: Retaining Meaning and Speed in High Dimensions. In *Advances in Neural Info. Proc. Systems 22*. 2009.
- [20] P. Ram, D. Lee, and A. G. Gray. Nearest-neighbor Search on a Time Budget via Max-Margin Trees. In *SIAM Intl. Conf. on Data Mining*, 2012.
- [21] J. H. Friedman, J. L. Bentley, and R. A. Finkel. An Algorithm for Finding Best Matches in Logarithmic Expected Time. *ACM Trans. Math. Softw.*, 1977.
- [22] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer, 1985.
- [23] J. McNames. A fast nearest-neighbor algorithm based on a principal axis search tree. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2001.
- [24] A.W. Moore. The Anchors Hierarchy: Using the triangle inequality to survive high dimensional data. In *Proc. of Conf. on Uncertainty in Artificial Intelligence*, 2000.
- [25] P. Ram, D. Lee, W. March, and A. Gray. Linear-time Algorithms for Pairwise Statistical Problems. In *Advances in Neural Info. Proc. Systems 22*. 2009.
- [26] A. Beygelzimer, S. Kakade, and J.C. Langford. Cover Trees for Nearest Neighbor. http://hunch.net/~jl/projects/cover_tree/paper/paper.pdf.
- [27] R. R. Curtin, J. R. Cline, N. P. Slagle, M. L. Amidon, and A. G. Gray. MLPACK: A Scalable C++ Machine Learning Library. In *BigLearning: Algorithms, Systems, and Tools for Learning at Scale*, 2011.
- [28] Y. LeCun. MNIST dataset, 2000. <http://yann.lecun.com/exdb/mnist/>.
- [29] C. L. Blake and C. J. Merz. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml/>, 1998.
- [30] J. Bennett and S. Lanning. The Netflix Prize. In *Proc. KDD Cup and Workshop*, 2007.
- [31] R. Lupton, J.E. Gunn, Z. Ivezic, G.R. Knapp, S. Kent, and N. Yasuda. The SDSS Imaging Pipelines. *Arxiv preprint astro-ph/0101420*, 2001.

- [32] S. Kim, F. Li, G. Lebanon, and I. Essa. Beyond Sentiment: The Manifold of Human Emotions. *Arxiv preprint arXiv:1202.1568*, 2011.
- [33] A. Torralba, R. Fergus, and W.T. Freeman. 80 Million Tiny Images: A large data set for nonparametric object and scene recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 2008.
- [34] L. Cayton. Accelerating nearest neighbor search on manycore systems. In *IEEE Intl. Parallel & Distributed Processing Symposium (IPDPS)*, 2012.
- [35] D. Lee, R. Vuduc, and A.G. Gray. A distributed kernel summation framework for general-dimension machine learning. In *SIAM Intl. Conf. on Data Mining*, 2012.